build anything together

Calculating what we can’t see: carbon emissions in the cloud

Leah Goldfarb &
Mary Thomas

18 November 2022 // Symfony Con 2022, France

platform.sh

Welcome, and thank you for joining us this morning, especially if you were able to enjoy the private access to the park last evening.

We are going to have an interesting discussion about calculating what we can’t see: that is, carbon emissions in the cloud.

// Mary Thomas

Manager, Data Analytics Platform.sh

- MS in Transportation Engineering
- Data systems + analytics expert
- Hiking + nature enthusiast

2

I am Mary Thomas, and I am Manager of Data Analytics at Platform.sh.

I originally have my master's degree in Transportation Systems Engineering, where I was exposed to the environmental policy and regulation with respect to the transport sector.

Now, I specialize in data systems and managing data platforms.

And, personally, I love hiking and being in nature, along with my two children.

And now, I will hand off to my colleague Leah.

// Leah Goldfarb

Environmental Impact Officer, Platform.sh

- PhD in physical chemistry: atmospheric focus
- Climate specialist
- Passionate about reducing tech's environmental impact
- Mom, cyclist, book-lover

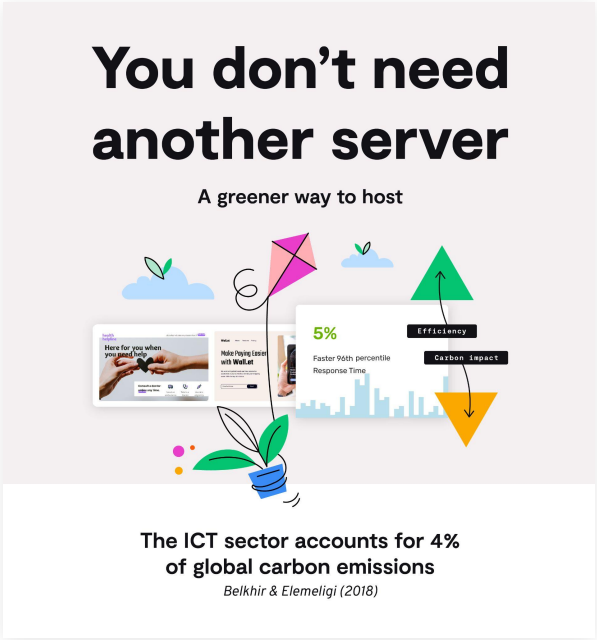

4

If you had a chance to join us for the keynote, you’ll know that you don’t need another server. Our solution - the shared responsibility model - accomplishes both performance improvement *and* lower carbon impact relative to traditional solutions. We can deploy to greener regions; we can achieve density and high compute resource utilization; we can reduce load times and identify bottlenecks.

// Outline

- **Climate projections**
- **What hyperscalers are doing**
- **Complexities of cloud carbon calculations**
  - \+ Learn and be transparent
- **A quick calculation**
- **What's next?**

// Net zero

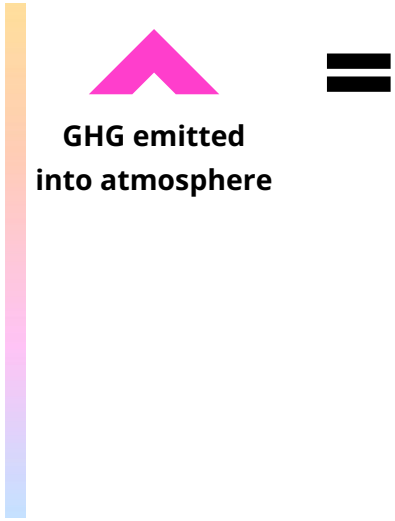

GHG emitted into atmosphere

GHG removed from atmosphere

// Net zero

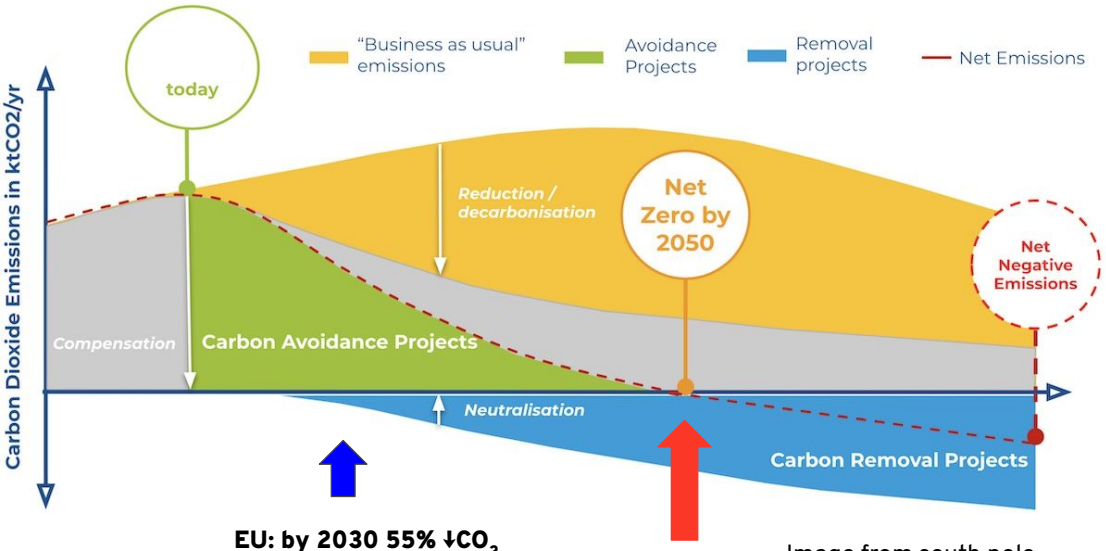

EU: by 2030 55% ↓CO_2

Image from south pole

// What needs to be done to meet the 2 °C Paris Agreement Goal

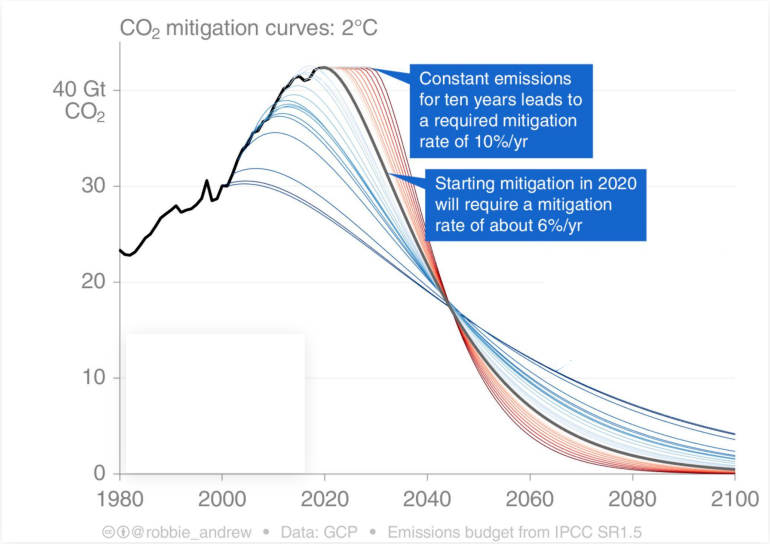

Climate Action Tracker

- Speaker: Leah
- Message(s)
  - + This is a curve showing a timeline until 2100 on the horizontal
  - + And Gt of CO2 on the vertical.
  - + The curves show the pathways to keep warming to 2 C by the end of of century, as layed out in the Paris agreement.
  - + The COP27 is scheduled to end today. Any while 1.5 ° C may be out of reach, 1.51 C should be our next goal.
  - + To do this we need to make massive cuts in our emissions to keep the devastating effects at bay
  - + So how do these cut start? We need to

+ reduce GHG emissions
+
+ EU 2015: cut 40% by 2030 (1990 is the ref year)
+ EU 2019: (GD): cut emissions by at least 55% by 2030

// Carbon in the Cloud*

***Location is important**

Mary starts here—----------

Carbon accounting is relatively new but it is becoming more common for large companies to report carbon audits,

We want to touch on some of the key concepts and important considerations for the cloud.

As you can imagine, at Platform.sh, FinOps (in essence, cloud financial and operational management) is an important practice, with our product that interacts with 5 different cloud providers. We recognized that the work we do to manage our resources directly relates to our environmental footprint. As our FinOps practice matures in the organization, so too can our environmental impact analysis.

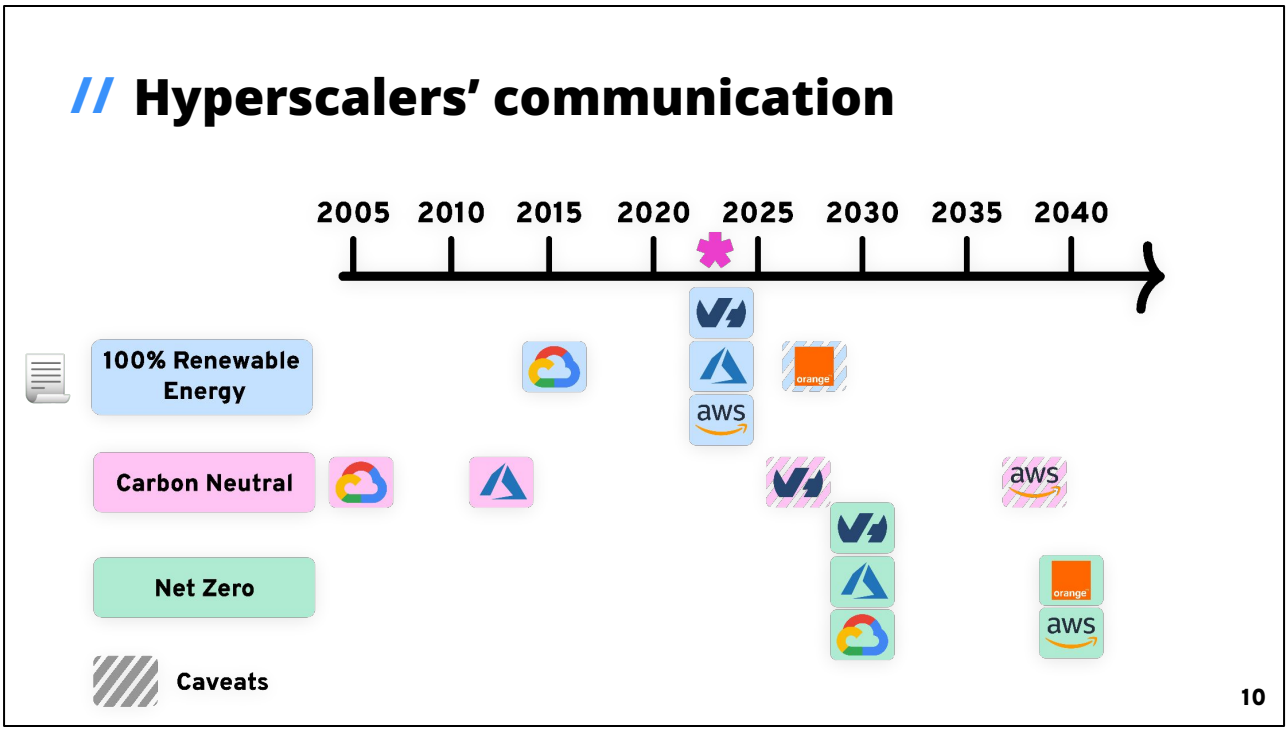

I am the messenger… here is what the providers have communicated.
Dotted lines means that there are caveats.

DESIGN: here's the svg for this diagram.

GCP: google cloud platform
Triangle :Microsoft Azure

AWS: Amazon Web Services

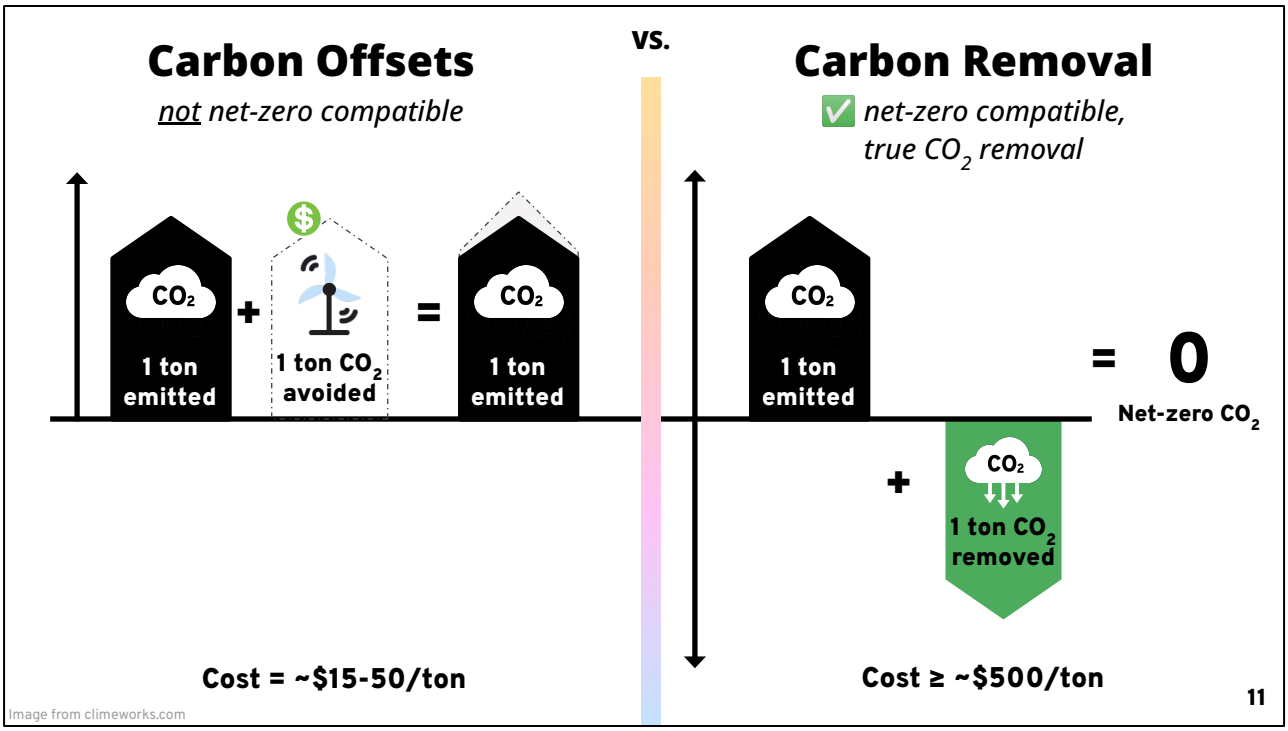

Could make a similar slide for PPAs

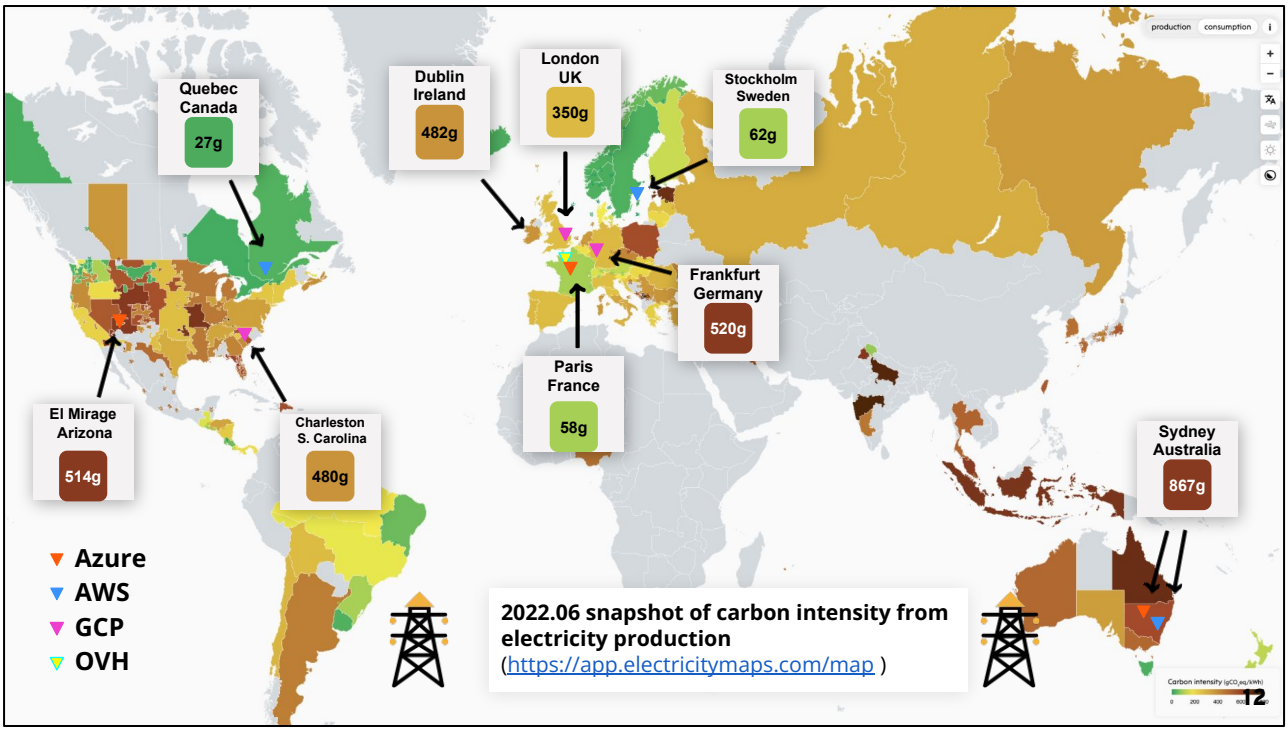

Gravelines is the site in the north of france

Based in the Denmark. Check some of their code on gittub.

CI may flux in time.
Numbers are from IEA.

Grid offerings, More on dedicated

// Market-based: not necessarily 1:1 carbon*

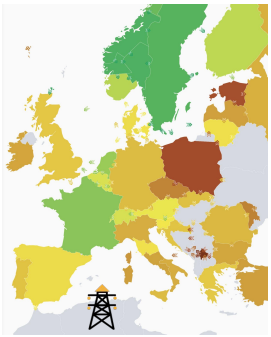

*Why your data centers should be located on clean electric grids.

Note, when PPAs are on the same elec. grid, market-based =location-based.

Market-based CO_2 dashboard: #

Market-based CO_2 dashboard: #

13

Missing element? Electricty

DESIGN: Here's the svg for the icon.

// Don't I have carbon info on my dashboard?

Sample AWS dashboard

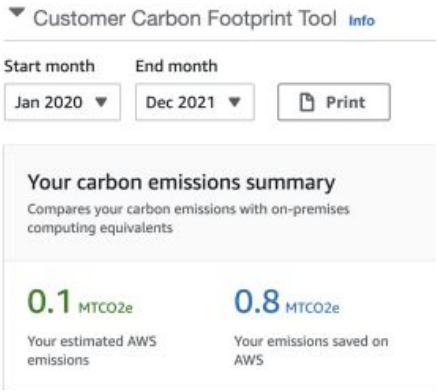

- **MS Azure, Amazon, ...dashboards can't be used in carbon audits**
  - \+ carbon auditor
  - \+ GCP dashboard shows location-based...
- **They often use a market-based approach**
  - \+ Even when location-based numbers are used
  - \+ There is something missing...

Missing element? Electricity
Metric Tons of Carbon
(see:https://aws.amazon.com/blogs/aws/new-customer-carbon-footprint-tool/)

// Carbon auditing can be complex

Mary starts here—----------

Carbon accounting is relatively new but it is becoming more common for large companies to report carbon audits,

We want to touch on some of the key concepts and important considerations for the cloud.

As you can imagine, at Platform.sh, FinOps (in essence, cloud financial and operational management) is an important practice, with our product that interacts with 5 different cloud providers. We recognized that the work we do to manage our resources directly relates to our environmental footprint. As our FinOps practice matures in the organization, so too can our environmental impact analysis.

But we know that, more so than FinOps, carbon auditing is quite

complex.

// The research is still emerging...

...but there is agreement that:

- **Cloud** is less resource intensive than on premises
- A **location-based approach** is the standard in carbon accounting
- It is a process: **transparency** is a necessity
- An accredited auditor should use **ghg protocol** for carbon accounting

...and we will continuing to refine our process as new peer-reviewed research is published

One complexity is that the research on GHG emissions in cloud computing is still emerging

but we have alignment on key points.

We know that cloud computing is less resource intensive than on premises networks.

We know that the standard for carbon accounting is a location-based approach, as Leah has just detailed.

We know that we must be transparent because this is an ever-improving process.

And we know that an accredited carbon audit will use the Greenhouse Gas Protocol standards.

Leah gave us an overview of the climate science that we have to consider when approaching the topic of GHG emissions at this

moment in time, but we have an obligation and commitment to continue refining our process as new research emerges from the scientific community.

// Scope 1, 2, 3 emissions

Defined by
GHG Protocol

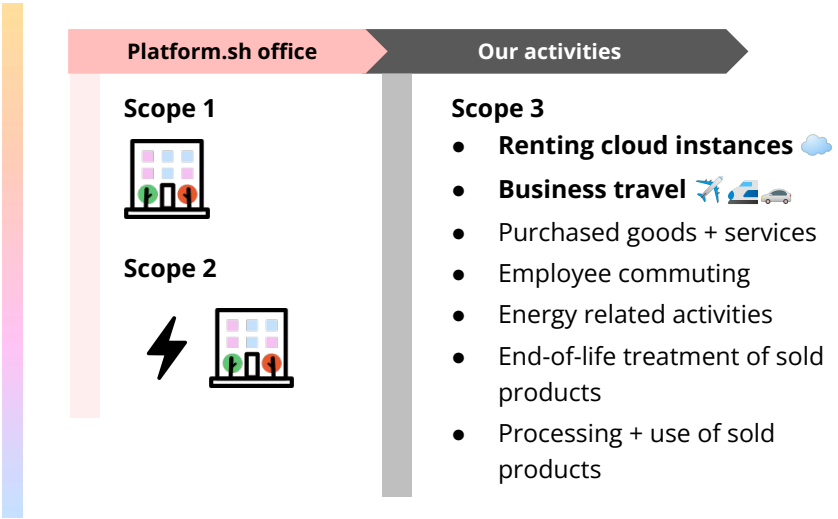

20

Another complexity that we have to consider is the way in which the GHG Protocol categorizes emissions, referred to as Scopes. The GHG protocol establishes comprehensive global standardized frameworks to measure and manage greenhouse gas emissions.

Scopes 1 and 2 are emissions that we own or control directly. Our Scope 1 is our Platform.sh Paris office, exclusive of electricity usage.
Scope 2 would be the electricity we purchase for the Paris office.

Scope 3 includes emissions that are a result of our activities - but emissions which we accumulate indirectly.
For example, business travel, like my and Leah's trains to get here; employees commuting to the office; etc.
Most pertinent to Platform.sh - what ends up being our biggest source of GHG emissions - is the emissions generated from running cloud instances / VMs.

GHG protocol was established by the WRI (World Resources Institute) and WBCSD (World Business Council for Sustainable Development)

// **At Platform.sh, we know that reducing our environmental impact benefits you**

Complexities aside, at Platform.sh, we know that reducing our environmental impact benefits you.

// Cloud methodology

Total CO_2eq emissions = **Embodied Emissions** + **Operational Emissions**

Embodied emissions

Small compared to operational emissions

- Estimated emissions from the manufacturing of data center servers

Operational emissions factors

- Capacity used: compute, storage, network
- Emissions factor for type of capacity
- Power usage efficiency
- Power mix

19

I want to be clear that, for transparency's sake, we use a 3rd-party auditor, Greenly, to complete our total CO2e emissions.

That said, what does a *typical* cloud carbon footprint calculation look like?

A very small portion of the total emissions relates to the manufacturing of the data center servers themselves.

The bulk of emissions are Operational, where we consider four variables.

// Operational emissions variables

Capacity used

vCPU, memory, and storage with a time dimension; networking without

Emission factor

Based on the type of the capacity use

Power usage effectiveness (PUE)

Represents the relationship between power used by servers and power used by cooling systems within a datacenter

Carbon intensity (CI)

Measures how “dirty” the electricity is for the grid where the datacenter exists

20

Our first variable to consider is the compute capacity used. We take into account vCPU, memory, and storage with a time dimension (for example, cpu-hours or GB-hours). Networking considers the transfer of data from one point to another without a time dimension.

Second, we look at the emission factor associated with the capacity used.

Third, we look at the Power Usage Effectiveness, which represents the relationship between power used by servers and power used by cooling systems within a datacenter. These are published as an average by each cloud provider, and remain stable over a 12-month period of time.

Finally, we look at the carbon intensity of the electric grid where the datacenter itself exists. Carbon intensity can change by the hour, and can vary significantly, by a factor of 2x or more.

// A VM* on AWS in France

Capacity used	Emission factor	Power usage effectiveness (PUE)	Carbon intensity (CI)		
					
35,040	× **1.844**	× **1.135**	× **0.0511**	÷ **1000**	
vCPU hours	**watts**	constant, reported by AWS	**kgCO₂eq / kWh**	W to kW	
1 year, 4 vCPU	for vCPU, based on 40% utilization		eu-west-3		

= **3.7 kg CO_2eq**

*** EC2 m5.xlarge**

cloudcarbonfootprint.org

In this example, we'll use methodology that is open source, and available at carbon cloud footprint dot org, which is a fantastic community, albeit a little simplistic and a bit behind current methodology. (another reason we use an independent auditing firm)

Source: https://www.cloudcarbonfootprint.org/docs/methodology

Emissions factor: (3.5-0.74)*.4 + 0.74 = 1.844

// We can reduce our carbon footprints together

These types of calculations can give us a clear picture of our carbon footprint, and with this information, we can work to reduce our carbon footprints together.

// What will Platform.sh do?

- Be **transparent** in our work
- Help customers use the **best region** for their projects
- Improve overall performance and **efficiency** of our orchestration layer
- Explore the potential to expose **location-based** carbon emissions data directly to customers
- Develop **new features** like dormant containers

27

What will Platform.sh

High on our list is to be transparent in all of our effort and communication around carbon accounting.

This year, we added the carbon intensity of each region in our console. Based on this information and other factors, we will help our customers use the best region for their projects.

Improving the performance and efficiency of our orchestration layer will reduce the environmental burden while supporting our customers' activity.

We know that the best practice is to use location-based methodology. Instead of relying on the market-based calculations provided by cloud providers, we are aiming to develop our own location based dashboard, and exploring what that might look like to expose the information directly to our customers.

And, of course, we will develop new features, like dormant containers, which will help put development containers to sleep automatically when not in use.

// What can you do?

- Make a conscious choice about the region in which you host
- Support your team/organization making environmental choices in development
- Optimize code to use resources efficiently

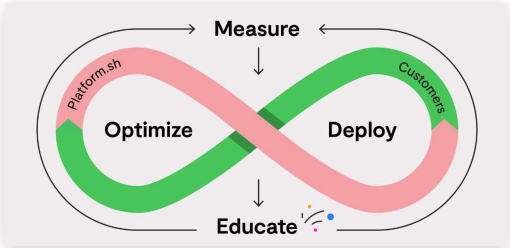

What can you developers do?

You can make a conscious decision about the region you choose to host your projects in.

You can support your team to make environmental choices in development - perhaps this means taking extra time for an architectural change that supports responsible code - it's organizational management.

And you can optimize your code using a tool like Blackfire - because ultimately you could be on the greenest of green servers and still be draining resources with inefficiencies.

It truly is a partnership between us and you.

// Looking forward

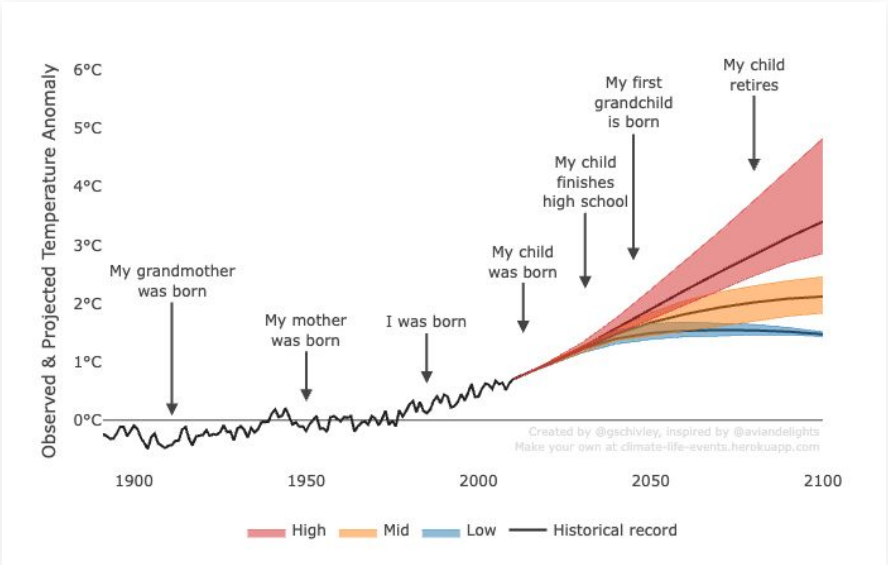

Inspired by [Sophie Lewis](#) IPCC AR5

We know a lot of these concepts may be new to a developer audience.

We know there are a lot of terms and methods, and research is constantly improving.

And Platform.sh sees it as our duty to be transparent about our carbon auditing and footprint, and to get to the point where we can help our customers do the same.

I hope you see the importance of considering the environmental impact of our collective cloud activity. This graph, inspired by the work and research of Sophie Lewis, shows major milestones generationally in my family, and it is interesting to see what kind of impact climate change could have by the time my child reaches my age. With collective action, aiming for the blue band, the future looks much brighter for the next generations to come.

AR5 - 2014

Thank you for your time and attention, Leah and I will be available at the Platform.sh booth for discussion and questions.